

# Risk Minimization in Structured Prediction using Orbit Loss

Danny Karmon  
Dept. of Computer Science  
Bar-Ilan University, Israel  
danny.karmon@biu.ac.il

Joseph Keshet  
Dept. of Computer Science  
Bar-Ilan University, Israel  
joseph.keshet@biu.ac.il

December 10, 2015

## Abstract

We introduce a new surrogate loss function called orbit loss in the structured prediction framework, which has good theoretical and practical advantages. While the orbit loss is not convex, it has a simple analytical gradient and a simple perceptron-like learning rule. We analyze the new loss theoretically and state a PAC-Bayesian generalization bound. We also prove that the new loss is consistent in the strong sense; namely, the risk achieved by the set of the trained parameters approaches the infimum risk achievable by any linear decoder over the given features. Methods that are aimed at risk minimization, such as the structured ramp loss, the structured probit loss and the direct loss minimization require at least two inference operations per training iteration. In this sense, the orbit loss is more efficient as it requires only one inference operation per training iteration, while yields similar performance. We conclude the paper with an empirical comparison of the proposed loss function to the structured hinge loss, the structured ramp loss, the structured probit loss and the direct loss minimization method on several benchmark datasets and tasks.

## 1 Introduction

There are three main differences between binary classification problems and structured prediction problems. First, the input to a binary classifier is a feature vector of a fixed length and the output is restricted to two possible labels, whereas in structured prediction both the input and the output are structured objects (a graph, an acoustic speech utterance, a sequence of words, an image). Second, the structured output space is potentially exponentially large (all possible phoneme or word sequences, all possible taxonomy graphs, all possible human poses, etc.). And third, while in binary classification the system's performance is evaluated using the error rate, i.e., 0-1 loss, in structured prediction each task often has its own evaluation metric or cost, such as word error rate in speech recognition, the BLEU score in machine translation, the NDCG score in information retrieval, or the intersection-over-union score in visual object segmentation. Some of these are involved functions, which are non-decomposable in the output space.

There is significant literature on learning parameters for structured prediction and graphical models. Ultimately, the goal in learning is to find the model parameters so as to minimize the expected cost, or risk, where the expectation is taken with respect to a random draw of input-output pairs from a fixed but unknown distribution. Since the expectation cannot be

evaluated because the underlying probability is unknown, and since the cost is often a non-convex combinatorial function (which is hard to minimize directly), the learning problem is formulated as an optimization problem where the parameters are found by minimizing a trade-off between a measure of the goodness of fit (loss) to the training data and a regularization term. In discriminative training, the loss function should be directly related to the cost between the model prediction and the target label, averaged over the training set.

The most common approaches to structured prediction, namely structured perceptron, structural support vector machine (SVM) and conditional random fields (CRF), do not directly minimize the risk. The structured perceptron (Collins, 2002) solves a feasibility problem, which is independent of the cost. In structural SVM (Joachims et al., 2005) the measure of goodness is a convex upper bound to the cost called structural hinge loss. It is based on a generalization of the binary SVM hinge loss to the structured case, and there is no guarantee for the risk. While there exists generalization bounds for the structured hinge loss (e.g., McAllester, 2006; Taskar et al., 2003), they all include terms which are not directly related to the cost, such as the Hamming loss, and inherently the structured hinge loss cannot be consistent as it fails to converge to the performance of the optimal linear predictor in the limit of infinite training data (McAllester, 2006). In CRFs the measure of goodness is the log loss function, which is independent of the cost (Lafferty et al., 2001). Smith and Eisner (2006) tried to address this shortcoming of CRFs and proposed to minimize the risk under the Gibbs measure. While it seems that this loss function is consistent, we are not aware of any formal analysis.

Recently, several works have focused on directly minimizing the expected cost. In particular, McAllester et al. (2010) presented a theorem stating that a certain perceptron-like learning rule, involving feature vectors derived from cost-augmented inference, directly corresponds to the gradient of the risk. Direct loss needs two inference operations per training iteration and is extremely sensitive to its hyper-parameter. Do et al. (2008) generalized the notion of the ramp loss from binary classification to structured prediction and proposed a loss function, which is a non-convex bound to the cost, and was found to be a tighter bound than the structured hinge loss function. The structured ramp loss also needs two inference operations per training iteration. Keshet et al. (2011) generalized the notion of the binary probit loss to the structured prediction case. The gradient of this non-convex loss function can be approximated by averaging over samples from the unit-variance isotropic normal distribution, where for each sample an inference with a perturbed weight vector is computed. In order to gain stability in the gradient computation, hundreds to thousands of inference operations are required per training iteration, hence the update rule is computationally-heavy.

The goal of this work is to propose a new learning update rule for structured prediction which results in fast training on one hand and aims at minimizing the risk on the other hand. We define a new loss function, called *orbit*, where its gradient has a close simple analytical form, which is very close to the structured perceptron update rule. We state a finite sample generalization bound for this loss function and show that it is consistent in the strong sense. That is, for any feature map (finite or infinite dimensional) the loss function yields predictors approaching the infimum risk achievable by any linear predictor over the given features. The update rule of this new loss involves one inference operation per training iteration, similar to the structured perceptron or the structural SVM, and hence faster (per training iteration) than ramp, probit and direct loss minimization. In a series of experiments we showed that the new loss function performs similar to other approaches that were designed to minimize the risk.

The paper is organized as follows. In Section 2 we state the problem formally. In Section 3 we introduce the new surrogate loss function and its update rule. In Section 4

we present the analysis for our new methods, including proofs for both consistency and generalization bound. In Section 5 we present a set of experiments and compare the new learning rule to other algorithms. We conclude the paper in Section 6.

## 2 Formal Settings

We formulate the structured supervised learning problem by setting  $\mathcal{X}$  to be an abstract set of all possible input objects and  $\mathcal{Y}$  to be an abstract set of all possible output targets. We assume that the input objects  $\mathbf{x} \in \mathcal{X}$  and the target labels  $\mathbf{y} \in \mathcal{Y}$  are drawn from an unknown joint distribution  $\rho$ . We define a set of  $d$  fixed mappings  $\phi : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$  called *feature functions* from the set of input objects and target labels to a real vector of length  $d$ .

Here we consider a linear decoder with parameters  $\mathbf{w} \in \mathbb{R}^d$ , such that the parameters weight the feature functions. We denote the score of label  $\mathbf{y}$  by  $\mathbf{w} \cdot \phi(\mathbf{x}, \mathbf{y})$ , given the input  $\mathbf{x}$ . The decoder predicts the label  $\hat{\mathbf{y}}_{\mathbf{w}}$  with the highest score:

$$\hat{\mathbf{y}}_{\mathbf{w}}(\mathbf{x}) = \operatorname{argmax}_{\mathbf{y} \in \mathcal{Y}} \mathbf{w} \cdot \phi(\mathbf{x}, \mathbf{y}) \quad (1)$$

Ideally, we would like to find the parameters  $\mathbf{w}$  that optimize the risk for unseen data. Formally, we define the *cost* function,  $\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})$ , to be a non-negative measure of error when predicting  $\hat{\mathbf{y}}_{\mathbf{w}}$  instead of  $\mathbf{y}$  as the label of  $\mathbf{x}$ . We assume that  $\ell(\mathbf{y}, \mathbf{y}) = 0$  for all  $\mathbf{y}$ . Often the desired evaluation metric is a utility function that needs to be maximized (like BLEU or NDCG) and then we define the cost to be 1 minus the evaluation metric.

Our goal is to minimize the risk:

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}(\mathbf{x}))]. \quad (2)$$

Since the distribution  $\rho$  is unknown, we use a training set  $\mathcal{S} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$  of  $m$  examples that are drawn i.i.d. from  $\rho$ , and replace the expectation in (2) with a mean over the training set and a regularization factor  $\frac{1}{2}\|\mathbf{w}\|^2$ . The cost is often a combinatorial non-convex quantity, which is hard to minimize, hence it is replaced with a *surrogate loss*, denoted  $\bar{\ell}(\mathbf{w}, \mathbf{x}, \mathbf{y})$ . Different algorithms use different surrogate loss functions. Overall the objective function of (2) transforms into the following objective function

$$\mathbf{w}^* = \operatorname{argmin}_{\mathbf{w}} \frac{1}{m} \sum_{i=1}^m \bar{\ell}(\mathbf{w}, \mathbf{x}_i, \mathbf{y}_i) + \frac{\lambda}{2} \|\mathbf{w}\|^2, \quad (3)$$

where  $\lambda$  is a trade-off parameter between the loss term and the regularization factor.

## 3 Orbit Loss

Denote by  $\Delta\phi(\mathbf{y}, \mathbf{y}')$  the difference between the feature functions of the labels  $\mathbf{y}, \mathbf{y}' \in \mathcal{Y}$ , respectively:

$$\Delta\phi(\mathbf{y}, \mathbf{y}') = \phi(\mathbf{x}, \mathbf{y}) - \phi(\mathbf{x}, \mathbf{y}').$$

Define by  $\delta\phi(\mathbf{y}, \mathbf{y}')$  the normalized version of  $\Delta\phi(\mathbf{y}, \mathbf{y}')$  as follows:

$$\delta\phi(\mathbf{y}, \mathbf{y}') = \begin{cases} \Delta\phi(\mathbf{y}, \mathbf{y}') / \|\Delta\phi(\mathbf{y}, \mathbf{y}')\| & \text{if } \mathbf{y} \neq \mathbf{y}' \\ \mathbf{0} & \text{if } \mathbf{y} = \mathbf{y}' \end{cases}. \quad (4)$$

The *orbit* surrogate loss function is formally defined as follows:

$$\bar{\ell}_{orbit}(\mathbf{w}, \mathbf{x}, \mathbf{y}) = \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)} \left[ \varepsilon > \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \right] \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}). \quad (5)$$

That is, the orbit loss is equal to the cost multiplied by the probability that the prediction score  $\mathbf{w} \cdot \phi(\mathbf{x}, \hat{\mathbf{y}}_{\mathbf{w}})$  plus a small number  $\varepsilon$  is greater than the score of the target label  $\mathbf{w} \cdot \phi(\mathbf{x}, \mathbf{y})$ .

We now derive the gradient-based learning rule for this loss function, which helps to describe some of its properties. The loss has a simple analytical gradient:

$$\nabla_{\mathbf{w}} \left[ \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)} [\varepsilon > \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}})] \ell(\mathbf{y}, \hat{\mathbf{y}}) \right] = \nabla_{\mathbf{w}} \left[ \frac{1}{\sqrt{2\pi}} \int_{\mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}})}^{\infty} e^{-z^2/2} dz \ell(\mathbf{y}, \hat{\mathbf{y}}) \right] \quad (6)$$

$$= -\frac{1}{\sqrt{2\pi}} e^{-|\mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}})|^2/2} \ell(\mathbf{y}, \hat{\mathbf{y}}) \delta\phi(\mathbf{y}, \hat{\mathbf{y}}). \quad (7)$$

The update rule of the orbit loss is the following:

$$\mathbf{w} \leftarrow (1 - \eta\lambda) \mathbf{w} + \eta e^{-|\mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})|^2/2} \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}). \quad (8)$$

Note that when the prediction label  $\hat{\mathbf{y}}_{\mathbf{w}}$  is close to the target label  $\mathbf{y}$  in terms of the decoding score, that is, when the term  $\mathbf{w} \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}}, \mathbf{y})$  is relatively small, the exponent is close to 1. Under this condition the update rule becomes

$$\mathbf{w} \leftarrow (1 - \eta\lambda) \mathbf{w} + \eta \ell(\mathbf{y}, \hat{\mathbf{y}}) \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}), \quad (9)$$

which generalizes the regularized structured perceptron's update rule (Collins, 2002; Zhang et al., 2014). Namely

$$\mathbf{w} \leftarrow (1 - \eta\lambda) \mathbf{w} + \eta \mathbb{1}\{\mathbf{y} \neq \hat{\mathbf{y}}_{\mathbf{w}}\} \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}), \quad (10)$$

where  $\mathbb{1}\{\pi\}$  is an indicator function, equals 1 if the predicate  $\pi$  holds and equals 0 otherwise.

A nice property of this update rule is that the cost function does not need to be decomposable in the size of the output. Decomposable cost functions are needed in order to solve the cost-augmented inference that is used in the training of structural SVMs (Joachims et al., 2005; Ranjbar et al., 2013), direct loss minimization (McAllester et al., 2010), or structured ramp loss (Do et al., 2008). It means that cost functions like word error rate or intersection-over-union can be used without being approximated.

Another property of the orbit loss is its similarity to the structured probit loss (Keshet et al., 2011). The probit loss was derived from the concept of stochastic decoder in the PAC-Bayesian framework (McAllester, 2003, 1998) and was shown to have both good theoretical properties and practical advantages (Keshet et al., 2011). The structured probit loss is defined as follows

$$\bar{\ell}_{probit}(\mathbf{w}, \mathbf{x}, \mathbf{y}) = \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}+\boldsymbol{\epsilon}})], \quad (11)$$

where  $\boldsymbol{\epsilon} \in \mathbb{R}^d$  is a  $d$ -dimensional isotropic Normal random vector. Note that the orbit loss (5) can be written as follows:

$$\begin{aligned} \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)} [\varepsilon > \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})] \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \\ = \mathbb{P}_{\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [-\boldsymbol{\epsilon} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) > \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})] \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}). \end{aligned} \quad (12)$$

The last equation holds since the inner product of an isotropic normal random vector  $\boldsymbol{\epsilon}$  with a unit-norm vector  $\delta\phi(\mathbf{y}, \hat{\mathbf{y}})$  is a zero-mean unit variance normal random variable. Writing the probability as the expectation of an indicator function, we have

$$\bar{\ell}_{orbit}(\mathbf{w}, \mathbf{x}, \mathbf{y}) = \mathbb{E}_{\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\mathbb{I}\{(\mathbf{w} + \boldsymbol{\epsilon}) \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) < 0\}] \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}). \quad (13)$$

Assuming  $\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}}$  for a value  $r$  small enough that  $\epsilon \in B(\mathbf{0}, r)$ , where  $B(\mathbf{0}, r)$  is a ball of radius  $r$  centered at  $\mathbf{0}$ , we can bring the cost function into the expectation term, that is

$$\mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\mathbb{I}\{(\mathbf{w} + \epsilon) \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}+\epsilon}) < 0\} \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}+\epsilon})] = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}+\epsilon})], \quad (14)$$

which is the structured probit loss.

## 4 Analysis

In this section we analyze the orbit loss. We derive a generalization bound based on the PAC-Bayesian theory, where we start by upper-bounding the probit loss with the orbit loss and then plugging it into a PAC-Bayesian generalization bound. Then we show that the decoder's parameters, which are estimated by optimizing the regularized orbit loss in the limit of infinite data, approach the infimum risk achievable by any linear decoder.

Recall that the structured probit loss is defined as:

$$\bar{\ell}_{probit}(\mathbf{w}, \mathbf{x}, \mathbf{y}) = \mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}+\epsilon})]. \quad (15)$$

The following theorem states a generalization bound for the probit loss function (Keshet et al., 2011).

**Theorem 1** (Generalization of probit loss). *For a fixed  $\gamma > 1/2$  we know that, with a probability of at least  $1 - \delta$  over the draw of the training data, the following holds simultaneously for all  $\mathbf{w}$ :*

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}(\mathbf{w}, \mathbf{x}, \mathbf{y})] \leq \frac{1}{1 - \frac{1}{2\gamma}} \left( \frac{1}{m} \sum_{i=1}^m \bar{\ell}_{probit}(\mathbf{w}, \mathbf{x}_i, \mathbf{y}_i) + \frac{\gamma}{2m} \|\mathbf{w}\|^2 + \frac{\gamma}{m} \ln \frac{1}{\delta} \right). \quad (16)$$

Later this generalization bound will help us state a similar bound for the orbit loss.

We now analyze the orbit loss. Let  $\eta$  be the minimal distance between the score of the predicted label  $\hat{\mathbf{y}}$  to the score of its closest *different* label  $\mathbf{y}'$  by a constant  $\eta$ :

$$\min_{\mathbf{y}' \neq \hat{\mathbf{y}}} \mathbf{w} \cdot \delta\phi(\hat{\mathbf{y}}, \mathbf{y}') \geq \eta \quad (17)$$

for  $\hat{\mathbf{y}} \neq \mathbf{y}'$ . The following lemma upper bounds the probit loss with the orbit loss.

**Lemma 2.** *For a finite  $\sigma > 0$  and a cost function  $\ell(\mathbf{y}, \mathbf{y}') \in [0, 1]$  for all  $\mathbf{y}, \mathbf{y}'$ , the following holds:*

$$\bar{\ell}_{probit}(\mathbf{w}/\sigma, \mathbf{x}, \mathbf{y}) \leq \bar{\ell}_{orbit}(\mathbf{w}/\sigma, \mathbf{x}, \mathbf{y}) + \sigma, \quad (18)$$

for  $\eta \geq \sigma \sqrt{2 \ln \frac{m}{\sigma}}$ .

For the brevity of the explanation we call  $\hat{\mathbf{y}}_{\mathbf{w}}$  the *predicted label* and we call  $\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}$  the *perturbed label*. The idea behind the proof is to split the structured probit loss cases of  $\epsilon$  for which the predicted label and the perturbed label are the same, and the case in which they differ. We show that the probability of the labels being equal is upper-bounded by the orbit loss, and the probability of the labels being different is upper-bounded by an exponential term that approaches zero when the norm of  $\epsilon$  approaches zero.

*Proof.* From the law of total expectation we have

$$\mathbb{E}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\ell(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \mathbf{y})] \leq \mathbb{E}_{\epsilon} [\mathbb{I}\{\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}}\} \ell(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \mathbf{y})] + \mathbb{P}_{\epsilon} [\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} \neq \hat{\mathbf{y}}_{\mathbf{w}}], \quad (19)$$

where we upper bound the cost by 1 in the second term.

First, let us focus on the first term of the inequality. For this term  $\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}}$ , which means that

$$\mathbb{E}_{\epsilon} \left[ \mathbb{1}\{\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}}\} \ell(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \mathbf{y}) \right] = \mathbb{P}_{\epsilon} \left[ \hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}} \right] \ell(\hat{\mathbf{y}}_{\mathbf{w}}, \mathbf{y}). \quad (20)$$

By definition of the inference rule (1) for any vector  $\mathbf{u}$ , we have  $\mathbf{u} \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{u}}, \mathbf{y}) \geq 0$  for all  $\mathbf{y}$ . Therefore the probability that  $\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}}$  can be expressed as follows:

$$\mathbb{P}_{\epsilon} \left[ \hat{\mathbf{y}}_{\mathbf{w}+\epsilon} = \hat{\mathbf{y}}_{\mathbf{w}} \right] = \mathbb{P}_{\epsilon} \left[ (\mathbf{w} + \epsilon) \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) \leq 0 \right] \quad (21)$$

which, in turn, can be expressed as

$$\mathbb{P}_{\epsilon} \left[ \mathbf{w} \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) \leq -\epsilon \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) \right] \leq \mathbb{P}_{\epsilon} \left[ \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \leq -\epsilon \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) \right], \quad (22)$$

where replacing  $\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}$  with  $\mathbf{y}$  increases the event size, thereby increasing the probability. Replacing the inner product of an isotropic normal random vector  $\epsilon$  with a unit-norm vector  $\delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})$  with a zero-mean unit variance normal random variable, we get:

$$\mathbb{P}_{\epsilon} \left[ \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \leq -\epsilon \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) \right] = \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)} \left[ \varepsilon > \mathbf{w} \cdot \delta\phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \right]. \quad (23)$$

The second term of the left-hand side of (19) can be expressed as follows:

$$\mathbb{P}_{\epsilon} [\hat{\mathbf{y}}_{\mathbf{w}+\epsilon} \neq \hat{\mathbf{y}}_{\mathbf{w}}] = \mathbb{P}_{\epsilon} [(\mathbf{w} + \epsilon) \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) > 0].$$

We have

$$\mathbb{P}_{\epsilon} [(\mathbf{w} + \epsilon) \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) > 0] = \mathbb{P}_{\epsilon} [\epsilon \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) > \mathbf{w} \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}}, \hat{\mathbf{y}}_{\mathbf{w}+\epsilon})] \quad (24)$$

$$\leq \mathbb{P}_{\epsilon} [\epsilon \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) > \eta]. \quad (25)$$

We finalized the proof by bounding the last equation for a  $\sigma$ -scaled version of  $\epsilon$ ,

$$\mathbb{P}_{\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\sigma \epsilon \cdot \delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\sigma\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}}) > \eta] = \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)} [\sigma \varepsilon > \eta] \leq \exp \left( -\frac{\eta^2}{2\sigma^2} \right) = \frac{\sigma}{m}, \quad (26)$$

where the first equation holds since the inner product of an isotropic normal random vector  $\epsilon$  with a unit-norm vector  $\delta\phi(\hat{\mathbf{y}}_{\mathbf{w}+\sigma\epsilon}, \hat{\mathbf{y}}_{\mathbf{w}})$  is a zero-mean unit variance normal random variable; and the second equation holds for  $\eta \geq \sigma \sqrt{2 \ln \frac{m}{\sigma}}$ . Using the union bound over the draw of a sample of size  $m$  concludes the proof.  $\square$

Plugging Lemma 2 into the bound of Theorem 1, we get the following generalization bound for the orbit loss.

**Theorem 3** (Generalization of orbit loss). *For a fixed  $\gamma > 1/2$  and assuming (17) holds with  $\eta \geq \sigma \sqrt{2 \ln(m/\sigma)}$ , we know that with a probability of at least  $1 - \delta$  over the draw of the training data the following holds true simultaneously for all  $\mathbf{w}$  and for all  $\sigma > 0$ :*

$$\begin{aligned} & \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}(\mathbf{w}/\sigma, \mathbf{x}, \mathbf{y})] \\ & \leq \frac{1}{1 - \frac{1}{2\gamma}} \left( \frac{1}{m} \sum_{i=1}^m \bar{\ell}_{orbit}(\mathbf{w}/\sigma, \mathbf{x}_i, \mathbf{y}_i) + \frac{\gamma}{2m\sigma^2} \|\mathbf{w}\|^2 + \sigma + \frac{\gamma}{m} \ln \frac{1}{\delta} \right). \end{aligned} \quad (27)$$

We will now prove that the orbit loss is consistent. We start with the observation that when the norm of the weight vector  $\mathbf{w}$  goes to infinity, the orbit loss approaches the cost:

**Lemma 4.**

$$\lim_{\alpha \rightarrow \infty} \bar{\ell}_{orbit}(\alpha \mathbf{w}, \mathbf{x}, \mathbf{y}) = \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}), \quad (28)$$

assuming that  $\ell(\mathbf{y}, \mathbf{y}) = 0$  for all  $\mathbf{y}$ .

*Proof.* Recall that  $\hat{\mathbf{y}}_{\mathbf{w}} = \arg \min_{\mathbf{y}'} \mathbf{w} \cdot \delta \phi(\mathbf{y}, \mathbf{y}')$  therefore  $\mathbf{w} \cdot \delta \phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \leq 0$ . Also note that scaling the parameters  $\mathbf{w}$  does not change the prediction,  $\hat{\mathbf{y}}_{\alpha \mathbf{w}} = \hat{\mathbf{y}}_{\mathbf{w}}$ . We have:

$$\begin{aligned} \lim_{\alpha \rightarrow \infty} \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)}[\varepsilon > \alpha \mathbf{w} \cdot \delta \phi(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})] \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) \\ = \mathbb{P}_{\varepsilon \sim \mathcal{N}(0,1)}[\varepsilon > -\infty] \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}) = \ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}). \end{aligned} \quad (29)$$

□

Consider the following training objective:

$$\hat{\mathbf{w}}_m = \arg \min_{\mathbf{w}} \frac{1}{m} \sum_{i=1}^m \bar{\ell}_{orbit}(\mathbf{w}, \mathbf{x}_i, \mathbf{y}_i) + \frac{\lambda_m}{2m} \|\mathbf{w}\|^2. \quad (30)$$

**Theorem 5** (Consistency of orbit loss). *For  $\hat{\mathbf{w}}_m$  defined by (30), if the sequence  $\lambda_m / \ln^2 m$  increases without bound, and the sequence  $\lambda_m / (m \ln m)$  converges to zero, then with a probability of one over the draw of the infinite sample we have:*

$$\begin{aligned} \lim_{m \rightarrow \infty} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}((\ln m) \hat{\mathbf{w}}_m, \mathbf{x}, \mathbf{y})] \\ = \inf_{\mathbf{w}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}}(\mathbf{x}))]. \end{aligned} \quad (31)$$

*Proof.* Set  $\delta = 1/m^2$ ,  $\sigma = 1/\ln m$ , and  $\gamma_m = \lambda_m / \ln^2 m$  into the bound (27). We decompose  $\mathbf{w}$  into a scalar  $\alpha$ , corresponding to the norm of  $\mathbf{w}$ , and a unit norm vector  $\mathbf{w}^*$ . Last, using Chernoff we upper-bound  $\frac{1}{m} \sum_{i=1}^m \bar{\ell}_{orbit}(\mathbf{w}, \mathbf{x}_i, \mathbf{y}_i)$  by  $\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{orbit}(\mathbf{w}, \mathbf{x}, \mathbf{y})] + \sqrt{\ln m / m}$  to get

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}((\ln m) \mathbf{w}_m, \mathbf{x}, \mathbf{y})] \leq \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}((\ln m) \alpha \mathbf{w}^*, \mathbf{x}, \mathbf{y})] \quad (32)$$

$$\begin{aligned} \leq \frac{1}{1 - \frac{\ln^2 m}{2\lambda_m}} \left( \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{orbit}(\alpha \mathbf{w}^*, \mathbf{x}, \mathbf{y})] \right. \\ \left. + \sqrt{\frac{\ln m}{m}} + \frac{\lambda_m \alpha^2}{2m} + \frac{1}{\ln m} + \frac{2\lambda_m}{m \ln m} \right), \end{aligned} \quad (33)$$

where  $\mathbf{w}_m$  is the minimizer of the right-hand side of the bound in (27), as well as of the optimization problem (30). Taking the limit when the number of examples  $m$  approaches infinity on both sides we have

$$\lim_{m \rightarrow \infty} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}((\ln m) \mathbf{w}_m, \mathbf{x}, \mathbf{y})] \leq \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{orbit}(\alpha \mathbf{w}^*, \mathbf{x}, \mathbf{y})] \quad (34)$$

Noting that by

$$\mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\bar{\ell}_{probit}(\mathbf{w}, \mathbf{x}, \mathbf{y})] \geq \inf_{\mathbf{w}} \mathbb{E}_{(\mathbf{x}, \mathbf{y}) \sim \rho} [\ell(\mathbf{y}, \hat{\mathbf{y}}_{\mathbf{w}})], \quad (35)$$

and letting  $\alpha$  approach infinity using Lemma 4 concludes the proof. □

## 5 Experiments

We evaluated the performance of the orbit loss by executing a number of experiments on several domains and tasks and compared the results with other approaches that are aimed at risk minimization, namely direct loss minimization (McAllester et al., 2010), structured ramp loss (Do et al., 2008), and structured probit loss (Keshet et al., 2011). For a reference, we present results for the structured perceptron, as we wanted to stress the empirical differences between the update rule in (9) and the one in (10), as well as for the structured hinge loss.

### 5.1 MNIST

In our first experiment we tested the orbit update rule on a multiclass problem. MNIST is a dataset of handwritten digit images (10 classes). It is divided into a training set of 50,000 examples, a test set of 10,000 examples and a validation set of 10,000 examples. We preprocess the data by normalizing the input images, and reducing the dimension from the original 784 attributes to 100 using PCA.

We used the orbit update rule as in (8). We defined the weight vector  $\mathbf{w}$  as a concatenation of 10 weight vectors  $\mathbf{w} = (\mathbf{w}^0, \mathbf{w}^1, \dots, \mathbf{w}^9)$ , each corresponding to one of the 10 digits. The update rule of example  $(\mathbf{x}_i, y_i)$ ,  $\mathbf{x}_i \in \mathbb{R}^{100}$ ,  $y_i \in \{0, \dots, 9\}$  can be simplified based on Kesler’s construction (Crammer and Singer, 2001) as follows:

$$\begin{aligned}\mathbf{w}^{y_i} &\leftarrow (1 - \eta\lambda) \mathbf{w}^{y_i} + \eta e^{-|\mathbf{w}^{y_i} \cdot \mathbf{x}_i - \mathbf{w}^{\hat{y}} \cdot \mathbf{x}_i|^2/2} \ell(\hat{y}, y_i) \mathbf{x}_i \\ \mathbf{w}^{\hat{y}} &\leftarrow (1 - \eta\lambda) \mathbf{w}^{\hat{y}} - \eta e^{-|\mathbf{w}^{y_i} \cdot \mathbf{x}_i - \mathbf{w}^{\hat{y}} \cdot \mathbf{x}_i|^2/2} \ell(\hat{y}, y_i) \mathbf{x}_i \\ \mathbf{w}^r &\leftarrow (1 - \eta\lambda) \mathbf{w}^r \quad \text{for all } r \neq y_i, \hat{y}\end{aligned}$$

Note that the exponent values throughout the training were very close to 1 and, practically, the update rule (9) could be used.

To properly evaluate the orbit loss we ran the experiment with two different cost functions for  $\ell(\hat{y}, y)$ : 0-1 loss and a semi-randomized matrix. We did so because the update rule (9) is identical to the structured perceptron update rule under the 0-1 loss.

In the first case, we set  $\eta = \eta_0/\sqrt{t}$ , for  $\eta_0 = 0.1$ ,  $t$  is the iteration number, and  $\lambda = 0.001$ . We also trained a multiclass perceptron with  $\eta = 1$ , and  $\lambda = 0$ , and a multiclass SVM with  $C = 0.01$  (Crammer and Singer, 2001). All of the hyper-parameters were chosen on the validation set. In all of the experiments we ran 4 epochs over the training data and used a linear kernel.

The results are given in Table 1 and suggest that there is a slight advantage for the orbit loss over the other algorithms. Recall that we previously showed that the perceptron is a special case of orbit loss and under this setting, in which  $\ell(\hat{y}, y) = 0$ -1 loss, hence the only difference between the results in the table is due to the regularization factor used with the orbit loss.

Table 1: Error rate of MNIST trained and evaluated with 0-1 loss.

Algorithm	Error rate (0-1 loss)
Multiclass perceptron	8.54%
Multiclass SVM	8.51%
Orbit	<b>8.34%</b>



As mentioned above, this experiment was executed once again, setting the cost function,  $\ell(\hat{y}, y)$ , to be a semi-randomized matrix. We generated a randomized cost matrix of size  $10 \times 10$ , such that the elements on the diagonal were all 0, and the rest of the elements were chosen uniformly at random to be either 1 or 2. We trained multiclass perceptron, multiclass SVM, and orbit using the following hinge loss for the cost function:

$$\bar{\ell}_{hinge}(\mathbf{w}, \mathbf{x}, y) = \max_{\hat{y}} \left[ \ell(y, \hat{y}) - \mathbf{w}^y \cdot \mathbf{x} + \mathbf{w}^{\hat{y}} \cdot \mathbf{x} \right] \quad (36)$$

To ensure reliability, we ran the second experiment for each algorithm with 10 different sampled matrices and averaged the results. The results are presented in Table 2. The results show a clear advantage for the orbit loss update rule in regards to the task loss. The reason is that the orbit loss can take advantage of minimizing a non 0-1 loss, as compared to perceptron.

Table 2: Error rate of MNIST trained and evaluated using the randomized evaluation metric. The orbit and SVM were trained using the randomized evaluation metric. The reported results are average over 10 randomized evaluation metrics.

Algorithm	Error rate (0-1 loss)	Cost
Multiclass perceptron	8.56%	13.28%
Multiclass SVM	8.52%	12.84%
Orbit	<b>8.31%</b>	<b>11.13%</b>

Table 3: Percentage of correctly positioned phoneme boundaries, given a predefined tolerance on the TIMIT core test set (192 examples). The first two lines (marked with \*) present algorithms that were trained on a training set of 1796 examples and a validation set of 400 examples. The rest of the lines correspond to a training set of 150 examples and a validation set of 100 examples.

	$\tau$ -alignment accuracy [%]				$\tau$ -insensitive loss
	$t \leq 10\text{ms}$	$t \leq 20\text{ms}$	$t \leq 30\text{ms}$	$t \leq 40\text{ms}$	
Brugnara et al. (1993)*	79.7	92.1	96.2	98.1	-
Keshet et al. (2007)*	75.3	88.9	94.4	97.1	-
Structural SVM	79.4	90.5	95.3	97.3	0.45
Structured ramp-loss	72.7	85.4	93.9	95.9	0.53
Direct loss minimization	81.0	90.6	94.5	96.8	0.47
Orbit	<b>84.4</b>	<b>93.5</b>	<b>97.0</b>	<b>98.3</b>	<b>0.31</b>

## 5.2 Phoneme alignment

Our next experiment focused on the phoneme alignment, which is used as a tool in developing speech recognition and text-to-speech systems. This is a structured prediction task — the input  $\mathbf{x}$  represents a speech utterance, and consists of a pair  $\mathbf{x} = (\mathbf{s}, \mathbf{p})$  of a sequence of acoustic feature vectors (mel-frequency cepstral coefficients),  $\mathbf{s} = (\mathbf{s}_1, \dots, \mathbf{s}_T)$ , where  $\mathbf{s}_t \in \mathbb{R}^d$ ,  $1 \leq t \leq T$ ; and a sequence of phonemes  $\mathbf{p} = (p_1, \dots, p_K)$ , where  $p_k \in \mathcal{P}$ ,

$1 \leq k \leq K$  is a phoneme symbol and  $\mathcal{P}$  is a finite set of phoneme symbols. The lengths  $K$  and  $T$  can differ for different inputs, although typically  $T$  is significantly larger than  $K$ . The goal is to generate an alignment between the two sequences in the input. The output  $\mathbf{y}$  is a sequence  $(y_1, \dots, y_K)$ , where  $1 \leq y_k \leq T$  is an integer giving the start frame in the acoustic sequence of the  $k$ -th phoneme in the phoneme sequence. Hence the  $k$ -th phoneme starts at frame  $y_k$  and ends at frame  $y_{k+1} - 1$ .

For this task we used the TIMIT speech corpus for which there are published benchmark results (Brugnara et al., 1993; Hosom, 2009; Keshet et al., 2007). We divided a portion of the TIMIT corpus (excluding the SA1 and SA2 utterances) into three disjoint parts containing 1500, 1796 and 400 utterances, respectively. The first part was used to train a phoneme frame-based classifier, which given the pair of speech frame and a phoneme, returns the level of confidence that the phoneme was uttered in that frame. The output classifier is then used along with other features as a seven dimensional feature map  $\phi(\mathbf{x}, \mathbf{y}) = \phi((\mathbf{s}, \mathbf{p}), \mathbf{y})$  as described in Keshet et al. (2007).

The seven dimensional weight vector  $\mathbf{w}$  was trained on the second set of 150 aligned utterances for  $\tau$ -insensitive loss

$$\ell(\mathbf{y}, \hat{\mathbf{y}}) = \frac{1}{|\mathbf{y}|} \max \{ |\mathbf{y}_k - \hat{\mathbf{y}}_k| - \tau, 0 \}, \quad (37)$$

with  $\tau = 10$  ms. This cost measures the average disagreement between all of the boundaries of the desired alignment sequence  $\mathbf{y}$  and the boundaries of predicted alignment sequence  $\hat{\mathbf{y}}$  where a disagreement of less than  $\tau$  is ignored.

We trained the system with the orbit update rule where  $\eta = 1.0/\sqrt{t}$  and  $\lambda = 0.2$ ; the structured perceptron update rule; the structural SVM optimized using stochastic gradient descent with  $C=5$  (Shalev-Shwartz et al., 2011); structured ramp-loss with  $\eta = 1.0/\sqrt{t}$ ,  $\lambda = 0.4$ ; and direct loss minimization algorithm with  $\epsilon = 1.1$  on a reduced training set of 150 examples (out of 1796) and a reduced validation set of 100 examples (out of 400). We were not able to train the system with the probit loss in a reasonable time.

The results are given in Table 3. The results in the first 4 columns should be read as the accuracy (in percentage) that the prediction was within  $\tau$ . The higher the better. The last column of the table is the actual loss computed by (37) - the smaller the better. In those results the orbit update rule outperforms other algorithms, and yields state-of-the-art results.

We would like to note that as in the MNIST experiment, the exponent values in the update rule were very close to 1 and, practically, the update rule (9) could be used.

### 5.3 Vowel duration

In the problem of vowel duration measurement we are provided with a speech signal which includes exactly one vowel preceded and followed by consonants (i.e., CVC). Our goal is to predict the vowel duration accurately. Precise measurement of vowel duration in a given context is needed in many phonological experiments, and currently is done manually (Heller and Goldrick, 2014).

The speech signal is represented as a sequence of acoustic features  $\mathbf{x} = (x_1, x_2, \dots, x_T)$  where each  $x_i$  ( $1 \leq i \leq T$ ) is a  $d$ -dimensional vector representing acoustic parameters, such as high and low energy, pitch, voicing, correlation coefficient, and so on (we extract  $d=22$  acoustic features every 5 msec). We denote the domain of the feature vectors by  $\mathcal{X} \subset \mathbb{R}^d$ . The length of the input signal varies from one signal to another, thus  $T$  is not fixed. We denote by  $\mathcal{X}^*$  the set of all finite length sequences over  $\mathcal{X}$ . In addition, we denote by  $t_b \in \mathcal{T}$  and  $t_e \in \mathcal{T}$  the vowel onset and offset times, respectively, where  $\mathcal{T} = \{1, \dots, T\}$ . For brevity

Table 4: The cost function (38) divided into the onset and offset terms, as learned by different approaches that are aimed at minimizing the cost.

	Onset			Offset			Run-time
	$\tau_b=0\text{ms}$	$\tau_b=10\text{ms}$	$\tau_b=20\text{ms}$	$\tau_e=0\text{ms}$	$\tau_e=20\text{ms}$	$\tau_e=35\text{ms}$	
Perceptron	1.394	1.040	0.757	4.390	3.929	2.643	24m 18s
Probit	5.215	5.671	6.603	4.832	6.331	5.283	1d 12h 47m
Ramp loss	6.507	6.501	5.46	8.573	8.084	7.028	45m 36s
Direct loss	1.426	1.068	0.933	<b>3.919</b>	<b>3.406</b>	2.529	45m 21s
Orbit loss	<b>1.308</b>	<b>0.893</b>	<b>0.607</b>	4.178	3.707	<b>2.531</b>	<b>24m 41s</b>

we set  $\mathbf{t} = (t_b, t_e)$ . The typical duration of an utterance is around 2 sec. There were  $n=116$  feature functions that described the typical duration of a vowel, the mean high energy before and after the vowel onset, and so on. The cost function we use is:

$$\ell(\hat{\mathbf{t}}, \mathbf{t}) = [|\hat{t}_b - t_b| - \tau_b]_+ + [|\hat{t}_e - t_e| - \tau_e]_+, \quad (38)$$

where  $[\pi]_+ = \max\{0, \pi\}$ , and  $\tau_b, \tau_e$  are pre-defined constants. The above function measures the absolute differences between the predicted and the manually annotated vowel onsets and offsets. Since the manual annotations are not exact, we allow a mistake of  $\tau_b$  and  $\tau_e$  frames at the vowel onset and offset respectively.

We trained the system using the orbit update rule with  $\eta = 1/\sqrt{t}$ ,  $\lambda = 0.001$ ; the structured perceptron update rule; structured ramp loss with  $\eta = 0.1/\sqrt{t}$ ,  $\lambda = 0.8$ ; probit loss with the expectation approximated by a mean of 100 random samples  $\eta = 0.001/\sqrt{t}$ ,  $\lambda = 0.005$ ; and direct loss minimization with  $\eta = 0.1/\sqrt{t}$  and  $\epsilon=-1.52$ . All of those hyper-parameters were chosen for the validation set. The results are presented in Table 4 for different values of  $\tau_b$  and  $\tau_e$  in the cost function. It can be seen that the orbit is close to the direct loss minimization (differences of a frame or two on average) and is better than other approaches. Also note that as describe earlier, the efficiency of the orbit loss is similar to the structured perceptron update and better than other approaches.

## 6 Discussion and Future Work

We introduced a new surrogate loss function that offers an efficient and effective learning rule. We gave a qualitative theoretical analysis presenting a PAC-Bayesian generalization bound and a consistency theorem. Despite the fact that the consistency property concerns the training performance when the number of training examples is big, the proposed loss function was shown to perform well on several tasks, even when the training set was of small or medium size.

In terms of theoretical properties, we think that the theoretical analysis can be improved, and in particular we would like to have a better upper-bound of the probit loss in terms of the orbit loss, as expressed in Lemma 2, which depends on the minimal distance between the predicted label and its closest neighbor label. Anyways, it is clear that when the norm of the weight vector becomes large relative to the norm of the noise, the inference with the weight vector and the inference with the perturbed weight vector – both lead to the same predicted label with a high probability.

This work is part of our research on surrogate loss functions in the structured prediction setting. We believe that in order to understand what are good loss functions, we have to

understand the interrelationship between them. While we showed some relation between the orbit loss, the Perceptron and the probit loss, we still think that more work should be done. We are especially interested in understanding the connection between the orbit, the probit, and the direct loss minimization approach.

## References

- Brugnara, F., Falavigna, D., and Omologo, M. (1993). Automatic segmentation and labeling of speech based on hidden markov models. *Speech Communication*, 12:357–370.
- Collins, M. (2002). Discriminative training methods for hidden markov models: Theory and experiments with perceptron algorithms. In *Conference on Empirical Methods in Natural Language Processing*.
- Crammer, K. and Singer, Y. (2001). On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2:265–292.
- Do, C., Le, Q., Teo, C.-H., Chapelle, O., and Smola, A. (2008). Tighter bounds for structured estimation. In *Advances in Neural Information Processing Systems (NIPS) 22*.
- Heller, J. R. and Goldrick, M. (2014). Grammatical constraints on phonological encoding in speech production. *Psychonomic bulletin & review*, 21(6):1576–1582.
- Hosom, J.-P. (2009). Speaker-independent phoneme alignment using transition-dependent states. *Speech Communication*, 51:352–368.
- Joachims, I., Tsochantaridis, T., Hofmann, T., and Altun, Y. (2005). Large margin methods for structured and interdependent output variables. *Journal of Machine Learning Research*, 6:1453–1484.
- Keshet, J., McAllester, D., and Hazan, T. (2011). PAC-Bayesian approach for minimization of phoneme error rate. In *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- Keshet, J., Shalev-Shwartz, S., Singer, Y., and Chazan, D. (2007). A large margin algorithm for speech and audio segmentation. *IEEE Trans. on Audio, Speech and Language Processing*, 15(8):2373–2382.
- Lafferty, J., McCallum, A., and Pereira, F. (2001). Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In *Proceedings of the Eighteenth International Conference on Machine Learning (ICML)*, pages 282–289.
- McAllester, D. (2003). Simplified PAC-Bayesian margin bounds. In *Proceedings of the Sixteenth Annual Conference on Computational Learning Theory*.
- McAllester, D. (2006). Generalization bounds and consistency for structured labeling. In Schölkopf, B., Smola, A. J., Taskar, B., and Vishwanathan, S., editors, *Predicting Structured Data*, pages 247–262. MIT Press.
- McAllester, D., Hazan, T., and Keshet, J. (2010). Direct loss minimization for structured prediction. In *Advances in Neural Information Processing Systems (NIPS) 24*.
- McAllester, D. A. (1998). Some pac-bayesian theorems. In *Proceedings of the Eleventh Annual Conference on Computational Learning Theory*.

- Ranjbar, M., Lan, T., Wang, Y., Robinovitch, S. N., Li, Z.-N., and Mori, G. (2013). Optimizing nondecomposable loss functions in structured prediction. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(4):911–924.
- Shalev-Shwartz, S., Singer, Y., Srebro, N., and Cotter, A. (2011). Pegasos: Primal estimated sub-gradient solver for svm. *Mathematical programming*, 127(1):3–30.
- Smith, D. A. and Eisner, J. (2006). Minimum risk annealing for training log-linear models. In *Proc. of the COLING/ACL*, pages 787–794.
- Taskar, B., Guestrin, C., and Koller, D. (2003). Max-margin markov networks. In *Advances in Neural Information Processing Systems (NIPS) 17*.
- Zhang, K., Fujian, P., Su, J., and Zhou, C. (2014). Regularized structured perceptron: A case study on chinese word segmentation, pos tagging and parsing. *The 14th Conference of the European Chapter of the Association for Computational Linguistics (EACL)*, page 164.